

The Rapid Analysis of Scattering from Periodic Dielectric Structures Using Accelerated Cartesian Expansions (ACE)

Andrew D. Baczewski,^{1,2*} Nicholas C. Miller,¹ and Balasubramaniam Shanker^{1,2}

¹*Department of Electrical and Computer Engineering, Michigan State University,
East Lansing, MI, 48825, USA*

²*Department of Physics and Astronomy, Michigan State University,
East Lansing, MI, 48825, USA*

**Corresponding author: baczewsk@msu.edu*

The analysis of fields in periodic dielectric structures arise in numerous applications of recent interest, ranging from photonic bandgap (PBG) structures and plasmonically active nanostructures to metamaterials. To achieve an accurate representation of the fields in these structures using numerical methods, dense spatial discretization is required. This, in turn, affects the cost of analysis, particularly for integral equation based methods, for which traditional iterative methods require $\mathcal{O}(N^2)$ operations, N being the number of spatial degrees of freedom. In this paper, we introduce a method for the rapid solution of volumetric electric field integral equations used in the analysis of doubly periodic dielectric structures. The crux of our method is the ACE algorithm, which is used to evaluate the requisite potentials in $\mathcal{O}(N)$ cost. Results are provided that corroborate our claims of acceleration without compromising accuracy, as well as the application of our method to a number of compelling photonics applications. *Manuscript submitted to JOSA A.*

1. Introduction

The scattering of light from subwavelength dielectric arrays is fundamental to a number of contemporary problems in optics. In photonic bandgap (PBG) structures, gaps in the photonic density of states arise due to a modification of the photonic dispersion in the vicinity of Bragg planes brought about by some underlying periodicity [1]. These structures have

been explored for a wide range of applications, from waveguiding structures [2] to experimental tests of cavity quantum electrodynamics [3]. Similarly, plasmonic structures exploit the unusual dielectric behavior of certain metals (namely, the nobles) at optical frequencies to focus electromagnetic energy at subwavelength scales. This focusing is typically the result of the excitation of surface plasmon-polaritons (SPP) modes, with dispersion relations that lie below the free space light-line [4]. One means of coupling light into these modes is the periodic modulation of the surface morphology [5]. This physical phenomenon has been exploited for a number of applications, ranging from extraordinary transmission [6] to bridging the gap between electronic and optical circuitry [7]. Finally, metamaterial structures have been studied extensively in the hopes of finding media that can be homogenized, albeit at specific wavelengths and angles of incidence, in such a way that they effect the behaviors of a structure with a negative, near zero, or otherwise anomalous index of refraction [8]. While the noble metals are also often used in optical metamaterials, the associated losses can be too severe, leading to the development of active metamaterials, wherein dye molecules are used as gain-inducing dopants [9].

In all of the aforementioned applications, intuition and supplemental simulation have been successfully applied to the design and analysis of countless structures. However, this approach is bound to fail as designs become increasingly complex. It is the goal of computational electromagnetics and optics to supplant much of the trial and error associated with the design of such complex structures with rigorous *in silico* modeling. Typical systems that have been previously analyzed using full-wave methods are relatively small in terms of the number of spatial degrees of freedom, ranging from 100s-1,000s. However, a recent proliferation of research in PBG, SPP, and metamaterial systems with fine subwavelength features, and/or highly lossy/dispersive materials lead to periodic structures wherein the unit cell requires 10,000+ degrees of freedom to reliably represent the underlying physics.

Integral equation (IE) methods for periodic scattering problems present a number of salient advantages over full-wave differential methods. IE methods give rise to solution fields that exactly satisfy proper boundary conditions through the use of an appropriate Green's function, are free of numerical dispersion, and only require the discretization of scattering bodies. This is not to say that these methods are without their disadvantages, in particular, these methods give rise to dense matrices that require $\mathcal{O}(N^2)$ resources in terms of number of operations and storage to achieve an iterative solution. Further, the evaluation of matrix elements by numerical quadrature is somewhat onerous, as the periodic Green's function must be evaluated at each pair of source and testing integration nodes. To this end, a considerable body of research has developed with the goal of mitigating these costs through the design and implementation of fast algorithms.

Some popular fast methods include the Fast Multipole Method (FMM) [10], the Adaptive

Integral Method (AIM) [11], and numerous variations that address adaptivity, or low/high frequency bottlenecks. These methods have become very popular for the analysis of non-periodic scattering problems. While the framework for adapting FMM to periodic problems in statics has been around since [12], it was not until the mid-90's that FMM was adapted to periodic wave propagation [13]. Even so, it was used to reduce matrix fill time, rather than accelerating the iterative solution process- this was not realized until the work of Otani and Nishimura in 2008 [14]. Aside from tree-based methods, AIM-based methods have been very naturally employed in the analysis of periodic problems using both conventional IE methods [15], as well as FE-BI [16]. Further, a number of interpolatory methods have been applied to these problems [17], including highly efficient GPU implementations [18]. In this work, we expand upon this growing body of research in fast algorithms for periodic systems, and provide details of the extension of the method of Accelerated Cartesian Expansions (ACE) a hierarchical, tree-based method similar to the FMM, to doubly periodic dielectric arrays.

The ACE algorithm has previously been applied to the efficient computation of potentials of the form $R^{-\nu}$ [19], Lienard-Wiechert potentials [20], diffusion, lossy wave, and Klein-Gordon potentials [21], and periodic Helmholtz, Yukawa, and Coulomb potentials [22]. Like FMM, it is based upon a hierarchical decomposition of the computational domain mapped onto an octree data structure, wherein a distinction between near and farfield source-observer aggregates is made, and an addition theorem is used to effect the interaction of all bodies with linear scaling. Whereas this addition theorem is based upon spherical harmonics in the FMM, ACE utilizes Cartesian harmonics and takes the form of a generalized Taylor expansion. In doing so, the salient features that develop are:

- *Totally linear scaling:* in terms of both computational cost and storage.
- *Nearly kernel independent framework:* only multipole-to-local operators depend upon the explicit form of the Green's function.
- *Exact up/down tree traversal:* error is rigorously independent of tree height.
- *Amenability to non-uniform discretization:* multiscale structures can be handled very naturally.
- *Excellent low frequency accuracy:* conventional FMM for Helmholtz problems must be augmented for electrically dense problems [23, 24]. ACE has been shown to be very much complimentary to FMM in this regard [25].

It is this latter feature that we are particularly interested in exploiting in adapting ACE to periodic structures. In our primary applications of interest, the unit cells are at most

$1 - 2\lambda$, as the technologically compelling physics arises due to energy coupled into low order Bloch-Floquet modes. In many applications, the unit cell may be as small as $\lambda/4 - \lambda/10$.

The principal contributions of this work are two-fold, (i) an extension of the ACE algorithm that enables the efficient analysis of electrically dense periodic structures, and (ii) applications of this technology to a set of challenging problems. In presenting this work we will use the following layout. In Section 2, we provide a formal mathematical statement of the problem. In Section 3, we give details of the ACE algorithm, including the details necessary for its extension to volumetric periodic problems. Section 4 presents results that affirm our claims of linear scaling, convergence of our method to arbitrary accuracy, and utility in accelerating the analysis of a number of exemplary structures. Finally, Section 5 summarizes the contributions of this paper and gives a brief outline of work to come.

2. Problem Statement

Consider a doubly periodic array of dielectric scatterers, Ω_D , embedded in \mathbb{R}^3 , characterized by the dielectric function $\varepsilon(\vec{r}, \omega)$. The array is excited by an incident planewave, $\vec{E}_{inc}(\vec{r}, t) = \vec{E}_0 e^{i(\omega t - \vec{k}_{inc} \cdot \vec{r})}$, that gives rise to a polarization of the dielectric. In turn, this polarization field radiates a scattered field, $\vec{E}_{scat}(\vec{r}, t)$, such that the total field is given by $\vec{E}(\vec{r}, t) = \vec{E}_{inc}(\vec{r}, t) + \vec{E}_{scat}(\vec{r}, t)$. Considering only linear media, all fields will be time-harmonic, and we henceforth suppress an implicit factor of $e^{i\omega t}$ and all time-dependence.

The array upon which the geometry is arranged is characterized by the following 2-lattice, \mathcal{L}_2 :

$$\mathcal{L}_2 = \{\vec{t}_{m,n} = m\vec{a}_1 + n\vec{a}_2 \mid m, n \in \mathbb{Z}\} \quad (1)$$

Here, \vec{a}_1 and \vec{a}_2 are the lattice vectors, describing the periodicity of our array. Throughout, we will assume a square lattice, i.e., $|\vec{a}_1| = |\vec{a}_2|$ and $\vec{a}_1 \cdot \vec{a}_2 = 0$, though we note that our method can be extended to more irregular lattices (i.e., rectangular or skewed) with minimal modification. For completeness, we define the reciprocal lattice associated with \mathcal{L}_2 as:

$$\mathcal{L}_2^* = \{\vec{k}_{m,n} = m\vec{b}_1 + n\vec{b}_2 \mid m, n \in \mathbb{Z}\} \quad (2)$$

Where $\vec{a}_i \cdot \vec{b}_j = 2\pi\delta_{ij}$, δ_{ij} being the Kronecker Delta. A simple illustration of our problem is provided in Fig. 1.

We seek to resolve the unknown, $\vec{E}_{scat}(\vec{r})$, from which we can also compute quantities such as the scattering parameters, e.g., reflection, transmission, and absorption spectra, associated with Ω_D . In doing so, we utilize the volumetric equivalence principle to replace Ω_D with equivalent sources, $\vec{J}_V(\vec{r})$, radiating into a homogeneous space [26]. We relate $\vec{J}_V(\vec{r})$

to the electric displacement, $\vec{\mathbf{D}}(\vec{\mathbf{r}}) = \varepsilon(\vec{\mathbf{r}}, \omega) \vec{\mathbf{E}}(\vec{\mathbf{r}})$, via the following relationship:

$$\vec{\mathbf{J}}_V(\vec{\mathbf{r}}) = j\omega\kappa(\vec{\mathbf{r}})\vec{\mathbf{D}}(\vec{\mathbf{r}}), \quad \text{where } \kappa(\vec{\mathbf{r}}) = \frac{\varepsilon(\vec{\mathbf{r}}, \omega) - \varepsilon_0}{\varepsilon(\vec{\mathbf{r}}, \omega)} \quad (3)$$

With this substitution, any discontinuities in the normal component of $\vec{\mathbf{J}}_V(\vec{\mathbf{r}})$ across material interfaces is due to $\kappa(\vec{\mathbf{r}})$ rather than $\vec{\mathbf{D}}(\vec{\mathbf{r}})$, which facilitates the definition of local vector basis functions [27].

By relating $\vec{\mathbf{E}}_{scat}(\vec{\mathbf{r}})$ to $\vec{\mathbf{J}}_V(\vec{\mathbf{r}})$, and enforcing the identity of the total field, we arrive at the following Volume Integral Equation (VIE):

$$\vec{\mathbf{E}}_{inc}(\vec{\mathbf{r}}) = \vec{\mathbf{E}}(\vec{\mathbf{r}}) - \vec{\mathbf{E}}_{scat}(\vec{\mathbf{r}}), \quad \forall \vec{\mathbf{r}} \in \Omega_D \quad (4a)$$

$$\vec{\mathbf{E}}_{inc}(\vec{\mathbf{r}}) = \vec{\mathbf{D}}(\vec{\mathbf{r}})/\varepsilon(\vec{\mathbf{r}}, \omega) - i\omega\mu_0 \int_{\Omega_D} d\vec{\mathbf{r}}' g(\vec{\mathbf{r}}, \vec{\mathbf{r}}') \kappa(\vec{\mathbf{r}}') \vec{\mathbf{D}}(\vec{\mathbf{r}}') - \frac{i}{\omega\varepsilon_0} \nabla \int_{\Omega_D} d\vec{\mathbf{r}}' g(\vec{\mathbf{r}}, \vec{\mathbf{r}}') \nabla' \cdot (\kappa(\vec{\mathbf{r}}') \vec{\mathbf{D}}(\vec{\mathbf{r}}')) \quad (4b)$$

Here, $g(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$ is the free space Green's function for the Helmholtz equation in 3D. By taking advantage of periodicity, we restrict our consideration to a single unit cell, $\Omega_0 \subset \mathbb{R}^3$, and replace $g(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$ with the quasiperiodic Green's function for the Helmholtz equation in 3D, $g_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$. Using $g_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$ a radiation boundary condition is enforced receding away from the array, and Bloch-Floquet boundary conditions are enforced in the plane of \mathcal{L}_2 , giving the following relationship between fields/currents in different unit cells:

$$\vec{\mathbf{E}}(\vec{\mathbf{r}} + \vec{\mathbf{t}}_{m,n}) = e^{i\vec{\mathbf{k}}_{inc} \cdot \vec{\mathbf{t}}_{m,n}} \vec{\mathbf{E}}(\vec{\mathbf{r}}) \quad (5)$$

Equation (4) can then be reformulated as:

$$\vec{\mathbf{E}}_{inc}(\vec{\mathbf{r}}) = \vec{\mathbf{D}}(\vec{\mathbf{r}}, \omega)/\varepsilon(\vec{\mathbf{r}}, \omega) - \int_{\Omega_D^*} d\vec{\mathbf{r}}' \overline{\overline{\mathbf{G}}}_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}') \cdot [\kappa(\vec{\mathbf{r}}') \vec{\mathbf{D}}(\vec{\mathbf{r}}')] , \quad \Omega_D^* = \text{supp}(\vec{\mathbf{J}}_V) \cap \Omega_0 \quad (6)$$

Here, $\overline{\overline{\mathbf{G}}}_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$ is the Quasiperiodic Electric Dyadic Green's function, constructed from the scalar Green's function, $g_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$, per the formalism in [28].

To render a finite system of equations, Ω_D^* is represented as a tetrahedral tessellation with N faces, each of which is assigned a linear vector/Schaubert-Wilton-Glisson (SWG) basis function [27]. Galerkin testing is subsequently applied [26], yielding a linear system of N

equations in N unknowns of the form:

$$Z_{\mu\nu}I_\nu = V_\mu \quad (7a)$$

$$I_\nu = c_\nu, \quad V_\mu = \langle \vec{\mathbf{f}}_\mu(\vec{\mathbf{r}}), \vec{\mathbf{E}}_{inc}(\vec{\mathbf{r}}) \rangle \quad (7b)$$

$$Z_{\mu\nu} = \langle \vec{\mathbf{f}}_\mu(\vec{\mathbf{r}}), \vec{\mathbf{f}}_\nu(\vec{\mathbf{r}})/\varepsilon(\vec{\mathbf{r}}, \omega) \rangle - \langle \vec{\mathbf{f}}_\mu(\vec{\mathbf{r}}), \int d\vec{\mathbf{r}}' \vec{\mathbf{G}}_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}') \cdot [\kappa(\vec{\mathbf{r}}') \vec{\mathbf{f}}_\nu(\vec{\mathbf{r}}')] \rangle \quad (7c)$$

$$\vec{\mathbf{D}}(\vec{\mathbf{r}}) = \sum_{i=1}^N c_i \vec{\mathbf{f}}_i(\vec{\mathbf{r}}), \quad \forall \vec{\mathbf{r}} \in \Omega_D^* \quad (7d)$$

The expansion coefficients for $\vec{\mathbf{D}}(\vec{\mathbf{r}})$, c_i , can be resolved using an iterative method [29] in $\mathcal{O}(N^2)$ operations. These methods require the generation of a minimal sequence of vectors in the range of $Z_{\mu\nu}$ from which an approximate solution to Eqn. (7a) can be constructed such that $\|Z_{\mu\nu}I_\nu - V_\mu\|_2 \leq \epsilon$, where ϵ is some designated tolerance for error. This $\mathcal{O}(N^2)$ cost can be reduced via the application of fast solvers, wherein the dominant cost of the iterative solution process is reduced to $\mathcal{O}(N \log_2^\alpha(N))$ for $\alpha \in [0, 1]$. In the case of the ACE algorithm, $\alpha = 0$, and $\mathcal{O}(N)$ scaling has been demonstrated in numerous applications. In the Section 3, we give details of how this is achieved for the class of volumetric, doubly periodic vector Helmholtz problems described above.

3. ACE Algorithm

The crux of the ACE algorithm is the Generalized Taylor Expansion (GTE) expressed in the framework of Cartesian tensors:

$$f(\vec{\mathbf{r}} - \vec{\mathbf{r}}') = \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} \vec{\mathbf{r}}'^{(n)} \cdot n \cdot \nabla^{(n)} f(\vec{\mathbf{r}}) \quad (8)$$

Here, $f(\vec{\mathbf{r}})$ is some smooth function, $\cdot n \cdot$ indicates an n -fold tensor contraction, and the superscript (n) indicates a tensor of rank n . The premise of this equation, and of the ACE method, is to utilize the GTE in constructing an addition theorem for the periodic Green's function, wherein the source ($\vec{\mathbf{r}}'$) and observer ($\vec{\mathbf{r}}$) domains are separated. As in other analytical acceleration methods, it is this separation that enables the construction of a hierarchical decomposition of the potential integrals (rendered discretely as matrix-vector products) that are repeatedly evaluated in the iterative solution of Eqn. (7). To use this principle in achieving a linear scaling algorithm, we need to first designate when it is that we will utilize this separation to effect a portion of a matrix-vector product (i.e., nearfield vs. farfield metric), and then construct a mathematical framework that describes the process (i.e., tree traversal) in which it is applied. To this end, we describe:

- Construction of a hierarchical decomposition of Ω_D^* and its utility in identifying

nearfield vs. farfield basis function pairs.

- Application of the tree traversal process to effect a matrix-vector product.

3.A. Tree Construction

In the ACE algorithm, the matrix-vector product in Eqn. (7a) is approximated as:

$$Z_{\mu\nu}I_\nu \approx Z_{\mu\nu}^{near}I_\nu + \mathcal{T}^{ACE}(I_\mu) \quad (9)$$

Here, $Z_{\mu\nu}^{near}$ is a sparse matrix with $\mathcal{O}(N)$ entries which describe only interactions between pairs of basis functions that are in some metric near. The remaining term, \mathcal{T}^{ACE} is some composition of operators that effects the remaining farfield contribution to the total in $\mathcal{O}(N)$ cost. The distinction between nearfield and farfield is made apparent by mapping the basis functions subordinate to Ω_D^* onto a hierarchical decomposition, i.e., a tree. This is accomplished by embedding the unit cell inside of a cubic box and recursively subdividing it into increasingly smaller cubic boxes until some desired density of basis functions per box, σ , is achieved. This entire hierarchy of boxes is stored and represented in terms of a regular octree data structure. Boxes are referred to in terms of their genealogy, e.g., the box subordinate to a box at the level directly above it is the child to the superordinate box's parent. This genealogy is used in constructing the nearfield/farfield dictum as follows - the basis functions subordinate to two boxes constitute a farfield pair if:

- The separation between the original boxes and their nearest periodic images is, in all cases, at least a box length.
- Among the parents of both the original boxes and their images, at least one pair are in each other's nearfield.

The tree structure, and an exemplary application of this dictum are illustrated in Fig. 2.

It is important to note that the designation of nearfield and farfield does not carry the same physical significance as it typically does, but is only used in analogy with conventional free space tree-based methods. In practice, we implement this augmented rule by simply adding 2 auxiliary levels to our tree to account for the 8 nearest images of the unit cell, which are addressed but not explicitly 'filled'. As we utilize Morton ordering, the boxes lying in the original unit cell lie inside of a contiguous address space and we can easily distinguish between image boxes and the original boxes, but still utilize the same list building routines as in a free space code. This augmentation has a negligible effect on the computational cost, and can be considered a step in pre-processing as the image boxes are not explicitly traversed.

3.B. Tree Traversal

Next, we discuss the application of \mathcal{T}^{ACE} . As the ACE algorithm is formulated in the language of Cartesian tensors, it is necessary to first provide notational details. $\mathbf{M}_\mu^{(n)}$ and $\mathbf{L}_\mu^{(n)}$ are deemed the multipole and local expansions. These quantities are 3-vectors, indexed by μ , each component of which is a totally symmetric Cartesian tensors of rank n . A rank n Cartesian tensor on \mathbb{R}^3 is a set of 3^n quantities indexed by x , y , and z in each rank. Such a tensor is totally symmetric if these quantities are left invariant under permutation of the indices across rank, and will thus consist of $n(n+1)/2$ unique quantities. We denote an n -fold contraction between such tensors as $\cdot n \cdot$, and any vector quantity such as ∇ or $\vec{\mathbf{r}}$ taken as an n -fold tensor product with itself is indicated as $\nabla^{(n)}$ or $\vec{\mathbf{r}}^{(n)}$, respectively. Further details concerning tensor notation, as well as proofs of the identities that will follow can be found in [19].

For simplicity, we restrict ourselves to a 3-level tree, in which there will be no up/down tree traversal, noting that the details of a multilevel implementation can be found in [19, 25]. We begin by considering two disjoint subdomains of Ω_0 , Ω_s and Ω_o with centroids $\vec{\mathbf{r}}_s^c$ and $\vec{\mathbf{r}}_o^c$, deemed the source and observer domains, respectively. These subdomains are spatially separated in such a way that they separately fall inside leaf boxes that satisfy the farfield criterion described in Section 3.A. The first step in tree traversal is the construction of ACE multipole expansion about $\vec{\mathbf{r}}_s^c$, commonly referred to in the literature as the ‘charge-to-multipole’ or C2M step. For an ACE expansion truncated beyond order P , this amounts to calculating the $\mathcal{O}(P^3)$ unique entries of the following set of totally symmetric Cartesian tensors:

$$\mathbf{M}_\mu^{(n)} = \sum_{i=1}^{N_s} \frac{(-1)^n}{n!} w_{\mu,i} (\vec{\mathbf{r}}_i - \vec{\mathbf{r}}_s^c)^{(n)}, \quad 0 \leq n \leq P \quad (10)$$

Here, N_s is the number of quadrature points used in discretizing the contribution of the source integrals in Eqn. (7) due to basis functions in Ω_s , and $w_{\mu,i}$ is the weight associated with the μ th vector component of the i th quadrature point.

The coefficients contained in the \mathbf{M}_μ tensors, combined with a knowledge of the Taylor coefficients of the Green’s function allow us to calculate the fields at any point in $\vec{\mathbf{r}}_o^c$, from a local expansion about $\vec{\mathbf{r}}_o^c$, defined as:

$$\mathbf{L}_\mu^{(n)} = \sum_{m=n}^P \frac{1}{n!} \mathbf{M}_\mu^{(m-n)} \cdot (m-n) \cdot \nabla^{(m)} g_{per}(|\vec{\mathbf{r}}_o^c - \vec{\mathbf{r}}_s^c|), \quad 0 \leq n \leq P \quad (11)$$

Here, the aforementioned Taylor coefficients are contained in the set of totally symmetric tensors, $\nabla^{(m)} g_{per}(|\vec{\mathbf{r}}_o^c - \vec{\mathbf{r}}_s^c|)$, $0 \leq m \leq P$, referred to as the multipole-to-local (M2L) translation operator. As the quasi-periodic Green’s function is an infinite sum, an efficient means

of calculating the elements of the translation operator is necessary. To this end, the Ewald representation of the quasi-periodic scalar Helmholtz Green's function is used to generate these coefficients, a more detailed discussion of which can be found in [22]. We will make clear how the dyadic character of the Green's function used in the VIE can be accounted for using the scalar Green's function in what follows.

The final step in computing the fields in Ω_o due to sources in Ω_s , $\vec{\mathbf{E}}_\mu^{so}(\vec{\mathbf{r}})$, is termed the local-to-observer (L2O) step. This essentially amounts to the evaluation of the last two terms on the right hand side of Eqn. (4b). These fields are evaluated using the following relationship:

$$\vec{\mathbf{E}}_\mu^{os}(\vec{\mathbf{r}}) = \left(\mathbf{I}^{(2)} + \frac{\nabla^{(2)}}{k^2} \right) \cdot \mathbf{1} \cdot \sum_{n=0}^P \mathbf{L}_\nu^{(n)} \cdot \mathbf{n} \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}_o^c)^{(n)}, \quad \vec{\mathbf{r}} \in \Omega_o \quad (12)$$

Here, $\mathbf{I}^{(2)}$ is the rank 2 identity tensor. As $\mathbf{L}_\mu^{(n)}$ is a constant and the source and observer domains are disjoint, $\nabla^{(2)}$ can be applied directly to the $(\vec{\mathbf{r}} - \vec{\mathbf{r}}_o^c)^{(n)}$ tensors. From these fields, the testing integrals in Eqn. (7) can be evaluated, completing the farfield contribution to the matrix-vector multiplication. It is worth noting, that we can alternatively evaluate the tested field by transferring the $\nabla^{(2)}$ onto the source and testing basis functions prior to tree traversal. In this case, it is necessary to traverse an additional tree to account for the static contribution to the field. In practice, the difference between these two approaches comes down to a trade off between speed and accuracy. Transferring the derivatives typically results in a slightly more accurate evaluation of the field, whereas evaluating the field using the dyadic relation in Eqn. (12) is much more efficient. Consequently, the results presented in this paper have been obtained using the latter strategy.

3.C. Computational Cost

The cost structure of this scheme is identical to the one presented in [22], with the exception of some minor details. In [22] it is demonstrated that this periodized version of ACE provides $\mathcal{O}(N)$ evaluation of scalar potentials in terms of both FLOPs and memory. There are two primary differences between the scalar potentials analyzed previously and the vector potentials presented in this paper. First, a separate tree must be traversed for each vector component of the field, which will increase the number of operations and memory required for a tree traversal by a factor of either 3 or 4 depending upon whether or not the static contribution to the potential is evaluated using a separate tree, as previously discussed. Second, the quantity N is no longer the number of degrees of freedom (i.e, basis functions in this case), but proportional to this quantity. Instead, it is determined by the source and testing quadrature rules used in discretizing the necessary MoM integrals. While this doesn't affect the $\mathcal{O}(N)$ scaling of the method, it does affect the optimal number of degrees of freedom per leaf box used in minimizing cost. In the results presented, the average number of basis

functions per leaf box is ~ 20 .

4. Results

Next, we present several results that:

- Validate the proposed acceleration scheme and demonstrate convergence.
- Demonstrate the applicability of our method to the analysis of a numerical of established results.
- Explore the utility of the algorithm in analyzing a number of interesting photonic structures.

Both our ACE accelerated and reference codes were written in Fortran. They were compiled using the Intel Fortran Compiler with $-r8 -O3$ flags (double precision and optimization), and run in serial on an Intel Xeon E5620 at 2.4 GHz with 24 GB of RAM running Linux OS at the High Performance Computing Center at Michigan State University. In generating all results, we utilize the a diagonal-preconditioned TFQMR iterative method with a tolerance of $< 0.1\%$ in the L_2 norm. This choice of norm is implicit throughout the remainder of this work. All integrals over tetrahedra are evaluated with a 5 point source, 1 point testing rule, whereas all integrals over triangles are evaluated with a 7 point source, 1 point testing rule. This is of course, excepting singular integrations, which are evaluated using an analytic singularity subtraction in conjunction with quadrature [30]. All periodic Green's function evaluations, in both Ewald and Floquet representations, are converged to an accuracy of $< 0.001\%$. All infinite series in the periodic ACE translation operators are evaluated to a similar accuracy. In all cases, we note that we can significantly reduce the runtime of our algorithm by reducing the order of our ACE expansions and/or the convergence criterion for infinite summations. Unless otherwise indicated, we have utilized 5 levels and $P = 7$ to guarantee an error of $\sim 0.01\%$, in line with the results presented in [22] and in this work. In most scattering-based applications, a lower error tolerance/order of ACE expansion can be used with minimal recourse to the computed farfield observables. Finally, all exciting fields are normally incident on the plane of periodicity, excepting the results in Fig. (5).

The first two results are obtained using a 'kernel code' from which error convergence and scaling can be extracted independent of the framework of an integral equation solver. This 'kernel code' evaluates the following convolution, both with and without ACE:

$$\Phi(\vec{\mathbf{r}}) = \int_{\Omega_0} d\vec{\mathbf{r}}' g_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}') f(\vec{\mathbf{r}}') \quad (13a)$$

$$f(\vec{\mathbf{r}}') = \sum_{i=1}^N w_i \delta(\vec{\mathbf{r}}' - \vec{\mathbf{r}}_i) \quad (13b)$$

Here, N is the number of co-located point sources/observers, and w_i and $\vec{\mathbf{r}}_i$ denote the weight and position associated with the i th source, respectively. We are primarily concerned with the farfield contribution to this potential, i.e., for a fixed point observer, $\vec{\mathbf{r}}_i$, we restrict the integration in Eqn. (13a) to include only contributions from sources that are in the farfield of the observer as defined by the ACE algorithm, independent of whether or not ACE is used to evaluate the convolution. We denote this partial contribution to $\Phi(\vec{\mathbf{r}})$ as $\Phi_{far}(\vec{\mathbf{r}})$.

In Fig. (3) error convergence with the order of the ACE expansion is demonstrated using the ‘kernel code’. Here, the error was calculated using the following formula for 9 different values of P :

$$\epsilon_{far} = \sqrt{\frac{\sum_{i=1}^N \|\Phi_{far}^{ACE}(\vec{\mathbf{r}}_i) - \Phi_{far}^{direct}(\vec{\mathbf{r}}_i)\|^2}{\sum_{i=1}^N \|\Phi_{far}^{direct}(\vec{\mathbf{r}}_i)\|^2}} \quad (14)$$

In the results presented, $N = 1,000$ point sources were distributed randomly throughout a unit cell, $\Omega_0 = [0, 1) \times [0, 1) \times [0, 1)$ (i.e., the unit cube) and the weights, w_i , were chosen from a uniform random distribution on $[-1, 1)$. The phasing in $g_{per}(\vec{\mathbf{r}}, \vec{\mathbf{r}}')$ is characterized by an incident field with $\vec{\mathbf{k}}_{inc} = 2\pi\vec{\mathbf{a}}_1 \times \vec{\mathbf{a}}_2$, i.e., it is normally incident on the periodic array and has $\lambda = 1$. The results presented in Fig. (3) indicate that the ACE algorithm can be tuned to an arbitrary level of accuracy by increasing the order at which the associated expansions are truncated. We have previously demonstrated the effect of changing incidence angle and wavelength on error for a fixed value of P [22], and note here, that it has been found that convergence is largely unaffected by incidence angle, but not wavelength. Relative to the results presented, as wavelength decreases, convergence in P is slower, whereas it is typically faster as wavelength is increased. The conditions for this numerical experiment were chosen to demonstrate convergence in something of a worst case scenario for the applications at hand - as most of the technologically compelling physics for PBGs, plasmonic nanostructures, and metamaterials arise when the unit cell is subwavelength. It is worth noting that ϵ_{far} is evaluated relative to $\Phi_{far}(\vec{\mathbf{r}})$, rather than $\Phi(\vec{\mathbf{r}})$. As the nearfield contribution to the $\Phi(\vec{\mathbf{r}})$ is evaluated exactly, the error with respect to the total potential is typically an order of magnitude smaller.

The second set of ‘kernel code’ results are given in Fig. 4. Here, scaling of the precomputation and tree traversal times with the number of points sources N are presented. Timings are presented for 5 different values of P , demonstrating that linear scaling can be achieved at arbitrary precision. In generating the results in this Figure, N sources are distributed over a unit cell with $|\vec{\mathbf{a}}_1| = |\vec{\mathbf{a}}_2| = 1$, with a maximum out-of-plane dimension of 1, and $k = 2\pi$.

The number of levels is increased at each value of N , from 3 to 6, maintaining an average density of unknowns per leaf box of 80. The precomputation timings in Fig. 4(a) include the time required for tree construction as well as the calculation of all unique translation operators that will be later applied during tree traversal. This is a one time cost, and figures are provided to demonstrate two things: i) precomputation scales weakly with the number of sources/number of levels and ii) precomputation is negligible on the time scale required for solution, as will be evident from subsequent results. The tree traversal timings in Fig. 4(b) include the time required for all 5 steps of the tree traversal project. These results are intended to demonstrate that linear scaling is achieved for varying levels of accuracy (i.e., different values of P). A linear regression yields a scaling of $N^{1.01}$ for all values of P except for $P = 1$, in which the scaling is $N^{1.09}$.

Our first result illustrating integration with a VIE solver is given in Fig. 5. In it, we demonstrate the accuracy of our method in reproducing power transmission through an infinite dielectric slab of width 20 nm as the angle of incidence and polarization is varied at a fixed wavelength, $\lambda = 400$ nm. A similar result is presented in [31], wherein a layered-medium formulation is compared against an analytic calculation. Here, we utilize a periodic volumetric formulation wherein the unit cell has an edge length of 80 nm. As is evident, our ACE accelerated code reliably reproduces the analytic solution at both TE and TM polarization, at arbitrarily shallow incidence, for both positive and negative relative permittivities.

The results in Fig. 6 and Fig. 7 are intended to illustrate i) agreement of our code with established results from the literature and ii) acceleration of our code relative to our own unaccelerated reference code. The structure in Fig. 6 is an electromagnetic bandgap structure with applications in microwave engineering. The reference data is taken from [32], wherein a hybrid finite element-boundary integral formulation is utilized. Our ACE accelerated code is in good agreement with the established result, and we report a $37\times$ speedup relative to our unaccelerated code. The result in Fig. 7 is taken from [33] wherein an analytic calculation is performed using Mie theory. We again find good agreement between ACE and the established result and report a $46\times$ speedup relative to our reference code. This structure has been used elsewhere in the literature as validation of a periodic FMM code, wherein the authors report 4 – 18 minutes per solution on an 8 processor platform [14]. As a surface integral equation formulation is utilized, as well as parallelism, it is difficult to draw a direct comparison between our results. Taking these differences into consideration, it is safe to at least claim that our timings are competitive with the state of the art.

Our next result is intended to demonstrate the utility of the ACE algorithm in studying models for novel experimental applications. It has been understood for many years that a number of the optical effects common in the wings of butterflies, such as their deep coloring or iridescence, arise due to periodic nanostructures that occur naturally in the scales covering

these wings, on top of any chemical coloring (i.e. pigment) that may exist [34–36]. The blue coloring of the *Morpho* butterfly is partially due to a photonic bandgap (PBG) in the blue portion of the visible spectrum. This PBG is supported by a periodic ribbing in the material covering the butterfly’s wings on the scale of visible wavelengths [34]. Previously, Huang, et al. [37] have performed an extensive analysis of the optical properties of not just naturally occurring wings, but coated and synthetic wings made from Al_2O_3 . We have constructed a simple model to replicate the blue structural coloring of these types of structures, albeit one that is both thinner and oriented differently than the one in [37]. The reflectance of this structure, as well as its dimensions are given in Fig. 8. Dielectric constants for Al_2O_3 were interpolated from [38] and the resultant mesh has $N = 10,782$ unknowns. It is worth noting that we have artificially enlarged the unit cell for this problem to ensure that it is cubic. This both simplifies tree construction and increases the number of degrees of freedom for demonstrative purpose, and is not, in principal, necessary. In this result, the unaccelerated time to solution was extrapolated based upon the nearfield matrix fill time for ACE alone, neglecting the time for iterative solution, so the speedup factor of $46\times$ is actually representative of a lower bound on the acceleration.

Our final result is a demonstration of capability in solving a very densely discretized structure with a highly dispersive dielectric response. For inspiration, we turn to [39] wherein a yellow-light fishnet metamaterial made of layered Al_2O_3 and silver on an ITO substrate is simulated and fabricated. Here, we model a structure with the same periodicity and nanostrip widths but without tapering and an unmodified dielectric function for silver. While we utilize the well-established Johnson and Christy permittivity for silver [40], the authors of [39] increase the collision rate to account for surface roughness, grain boundaries, and size effects, thus we do not compare our results directly. Our resultant mesh has $N = 147,374$ unknowns, and was solved using ACE with 6 levels and $P = 5$. TFQMR was applied with diagonal preconditioning to arrive at an error of $< 3\%$ with an average of 188 iterations per frequency at 20 frequencies. The average time per iteration was 3.18 minutes, and the average total time to solution was 896 minutes. Extrapolating our nearfield matrix fill time indicates that the matrix fill process alone would require 256,500 minutes using our reference code with the same parameters, yielding a lower bound on the speedup of $\sim 286\times$.

5. Conclusion

We have demonstrated the extension of the ACE algorithm to the efficient analysis of volumetric integral equation-based scattering problems on periodic domains. Results were presented that indicate convergence of the kernel to arbitrary precision, as well as linear scaling in the evaluation of matrix-vector products. Corroborating results affirming the utility of our method in solving problems from the literature were also provided, as well as a number of

large problems that demonstrate capability in solving problems posed on densely discretized geometries pertinent to a number of subfields of the optics community. Work is presently under way in adapting our algorithm to both MPI and CUDA parallelism, surface formulations for dielectrics using the Generalized Method of Moments [41], and time domain problems.

Acknowledgments

The authors thank the National Science Foundation Graduate Research Fellowship (ADB) and the Michigan State University Summer Research Experience for Undergraduates (NCM) programs for supporting their research. More generally, the authors would like to acknowledge the body of work funded by NSF CCF-0729157 and NSF DMS-0811197 that have also made this work possible, as well as conversations and codes from Dan Dault, He Huang, Naveen Nair, and Melapudi Vikram. Finally, we thank the High Performance Computing Center (HPCC) at Michigan State University for access to computational resources.

References

1. S. John, “Strong localization of photons in certain disordered dielectric superlattices,” *Physical Review Letters* **58**, 2486–2489 (1987).
2. S.-Y. Lin, E. Chow, V. Hietala, P. Villeneuve, and J. Joannopoulos, “Experimental demonstration of guiding and bending of electromagnetic waves in a photonic crystal,” *Science* **282**, 274–276 (1998).
3. T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. Gibbs, G. Rupper, C. Ell, O. Shchekin, and D. Deppe, “Vacuum rabi splitting with a single quantum dot in a photonic crystal nanocavity,” *Nature* **432**, 200–203 (2004).
4. A. Otto, “Excitation of nonradiative surface plasma waves in silver by the method of frustrated total reflection,” *Zeitschrift fur Physik* **216**, 398–410 (1968).
5. W. Barnes, A. Dereuk, and T. Ebbesen, “Surface plasmon subwavelength optics,” *Nature* **424**, 824–830 (2003).
6. T. Ebbesen, H. Lezec, H. Ghaemi, T. Thio, and P. Wolff, “Extraordinary optical transmission through sub-wavelength hole arrays,” *Nature* **391**, 667–669 (1998).
7. E. Ozbay, “Plasmonics: Merging photonics and electronics at nanoscale dimensions,” *Science* **311**, 189–193 (2006).
8. V. Shalaev, “Optical negative-index metamaterials,” *Nature Photonics* **1**, 41–48 (2007).
9. S. Xiao, V. Drachev, A. Kildishev, X. Ni, U. Chettiar, H.-K. Yuan, and V. Shalaev, “Loss-free and active optical negative-index metamaterials,” *Nature* **466**, 735–738 (2010).
10. J. Song and W. Chew, “Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering,” *Microwave and Optical Technology Letters* **10**, 14–19 (1995).

11. E. Bleszynski, M. Bleszynski, and T. Jaroszewicz, "Aim: Adaptive integral method for solving large-scale electromagnetic scattering and radiation problems," *Radio Science* **31**, 1225–1251 (1996).
12. L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *Journal of Computational Physics* **73**, 325–348 (1987).
13. V. Rokhlin and S. Wandzura, "The fast multipole method for periodic structures," in "Antennas and Propagation Society International Symposium," (1994).
14. Y. Otani and N. Nishimura, "A periodic fmm for maxwell's equations in 3d and its applications to problems related to photonic crystals," *Journ* **227**, 4630–4652 (2008).
15. E. Bleszynski, M. Bleszynski, and T. Jaroszewicz, *Ultra-wideband, Short Pulse Electromagnetics 9* (Springer New York, 2010), chap. Rigorous Modeling of Electromagnetic Wave Interactions with Large Dense Systems of Descrete Scatterers, pp. 65–77.
16. T. Eibert and J. Volakis, "Adaptive integral method for hybrid fe/bi modelling of 3-d doubly periodic structures," *IEE Proceedings - Microwaves, Antennas and Propagation* **146**, 17–22 (1999).
17. S. Li, D. Van Orden, and V. Lomakin, "Fast periodic interpolation method for periodic unit cell problems," *IEEE Transactions on Antennas and Propagation* **58**, 4005–4014 (2010).
18. S. Li, B. Livshitz, and V. Lomakin, "Fast evaluation of helmholtz potential on graphics processing units (gpu)," *Journal of Computational Physics* **229**, 8463–8483 (2010).
19. B. Shanker and H. Huang, "Accelerated cartesian expansions- a fast method for computing of potential of the form $r^{-\nu}$ for all real ν ," *Journal of Computational Physics* **226**, 732–753 (2007).
20. M. Vikram and B. Shanker, "Fast evaluation of time domain fields in sub-wavelength source/observer distributions using accelerated cartesian expansions (ace)," *Journal of Computational Physics* **227**, 1007–1023 (2007).
21. M. Vikram, A. Baczewski, B. Shanker, and L. Kempel, "Accelerated Cartesian Expansion (ACE) based framework for the rapid evaluation of diffusion, lossy wave, and klein-gordon potentials," *Journal of Computational Physics* **229**, 9119–9134 (2010).
22. A. Baczewski and B. Shanker, "An $\mathcal{O}(n)$ method for rapidly computing periodic potentials using Accelerated Cartesian Expansions," submitted to *Journal of Computational Physics*, preprint available at arXiv:1107.3069 (2011). [Http://arxiv.org/abs/1107.3069v1](http://arxiv.org/abs/1107.3069v1).
23. L. Greengard, J. Huang, V. Rokhlin, and S. Wandzura, "Accelerating fast multipole methods for the helmholtz equation at low frequencies," *IEEE Computational Science and Engineering* **5**, 32–38 (1998).
24. L. Jiang and W. Chew, "A mixed-form fast multipole algorithm," *IEEE Transactions on*

- Antennas and Propagation **53**, 4145–4156 (2005).
25. M. Vikram, H. Huang, B. Shanker, and T. Van, “A novel wideband fmm for fast integral equation solution of multiscale problems in electromagnetics,” *IEEE Transactions on Antennas and Propagation* **57**, 2094–2104 (2009).
 26. A. Peterson, S. Ray, and R. Mittra, *Computational Methods for Electromagnetics* (Wiley-IEEE Press, 1998).
 27. D. Schaubert, D. Wilton, and A. Glisson, “A tetrahedral modeling method for electromagnetic scattering by arbitrarily shaped inhomogeneous dielectric bodies,” *IEEE Transactions on Antennas and Propagation* **32**, 77–85 (1984).
 28. R. Harrington, *Time-Harmonic Electromagnetic Fields* (Wiley-IEEE Press, 2001).
 29. Y. Saad, *Iterative Methods for Sparse Linear Systems* (SIAM, 20003).
 30. D. Wilton, S. Rao, A. Glisson, D. Schaubert, O. Al-Bundak, and C. Butler, “Potential integrals for uniform and linear source distributions on polygonal and polyhedral domains,” *IEEE Transactions on Antennas and Propagation* **32**, 276–281 (1984).
 31. G. Kobidze, B. Shanker, and D. Nyquist, “Efficient integral-equation-based method for accurate analysis of scattering from periodically arranged nanostructures,” *Physical Review E* **72** (2005).
 32. T. Eibert, J. Volakis, D. Wilton, and D. Jackson, “Hybrid fe/bi modeling of 3-d doubly periodic structures utilizing triangular prismatic elements and an mpie formulation accelerated by the ewald transformation,” *IEEE Transactions on Antennas and Propagation* **47**, 843–850 (1999).
 33. M. Inoue, K. Ohtaka, and S. Yanagawa, “Light scattering from macroscopic spherical bodies. ii. reflectivity of light and electromagnetic localized state in a periodic monolayer of dielectric spheres,” *Physical Review B* **25** (1982).
 34. M. Srinivasarao, “Nano-optics in the biological world: Beetles, butterflies, birds, and moths,” *Chem. Rev.* **99**, 1935–1961 (1999).
 35. P. Vukusic and J. Sambles, “Photonic structures in biology,” *Nature* **424**, 852–855 (2003).
 36. A. Parker and H. Townley, “Biomimetics of photonic nanostructures,” *Nature Nanotechnology* **2**, 347–353 (2007).
 37. J. Huang, X. Wang, and Z. Wang, “Controlled replication of butterfly wings for achieving tunable photonic properties,” *Nano Letters* **6**, 2325–2331 (2006).
 38. F. Gervais, *Handbook of Optical Constants of Solids* (Academic Press, Inc., 1991), vol. 2, pp. 761–775.
 39. S. Xiao, U. Chettiar, A. Kildishev, V. Drachev, and V. Shalaev, “Yellow-light negative-index metamaterials,” *Optics Letters* **34**, 3478–3480 (2009).
 40. P. Johnson and R. Christy, “Optical constants of the noble metals,” *Physical Review B* **6**, 4370–4379 (1972).

41. N. Nair and B. Shanker, “Generalized method of moments: a framework for analyzing scattering from homogeneous dielectric bodies,” *Journal of the Optical Society of America A* **28**, 328–340 (2011).




Fig. 1. Illustration of the periodic scattering problem described in Section 2.




Fig. 2. Top-down view of the geometry illustrated in Fig. 1 with a 4 level octree structure superimposed. Interaction lists are indicated for the dark blue box. Boxes highlighted in blue are in the nearfield, whereas light blue boxes are in the farfield. The interaction between sources in the dark blue box and boxes highlighted in red are effected at a higher level.






Fig. 3. Error convergence for $\Phi_{far}(\vec{r})$ evaluated using Eqn. (14).



(a) Scaling of precomputation time with the number of point sources for expansions truncated at different values of P



(b) Scaling of tree traversal time with the number of point sources for expansions truncated at different values of P .

Fig. 4. Timing results for ACE 'kernel code'.




Fig. 5. Validation of our ACE accelerated code against an analytic solution for scattering from a homogeneous dielectric slab of width 20 nm with $\epsilon_r = \pm 4$.




Fig. 6. Validation against scattering from an electromagnetic bandgap (EBG) structure solved using FE-BI in [32]. Discretization has $N = 6,030$ unknowns. Average time to solution (without acceleration): ~ 813 minutes. Average time to solution (ACE acceleration): ~ 23 minutes. Total speedup: $\sim 37\times$.

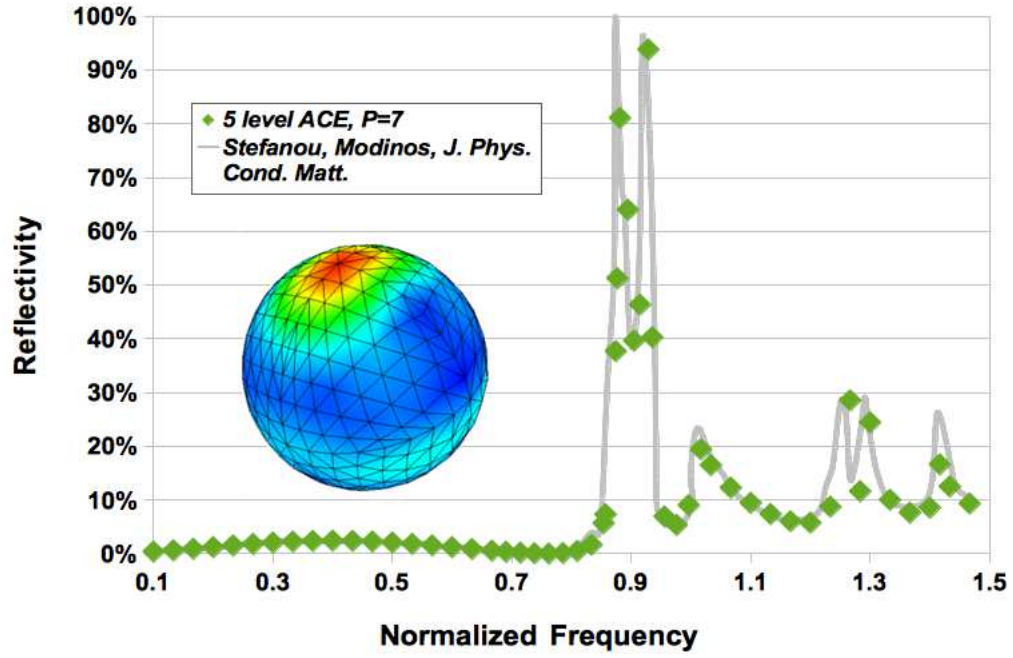


Fig. 7. Validation against scattering from an array of polystyrene spheres ($\varepsilon_r = 2.56$) solved analytically in [33]. Discretization has $N = 7,328$. Average time to solution (without acceleration): ~ 2006 minutes. Average time to solution (ACE acceleration): ~ 43 minutes. Total speedup: $\sim 46\times$.

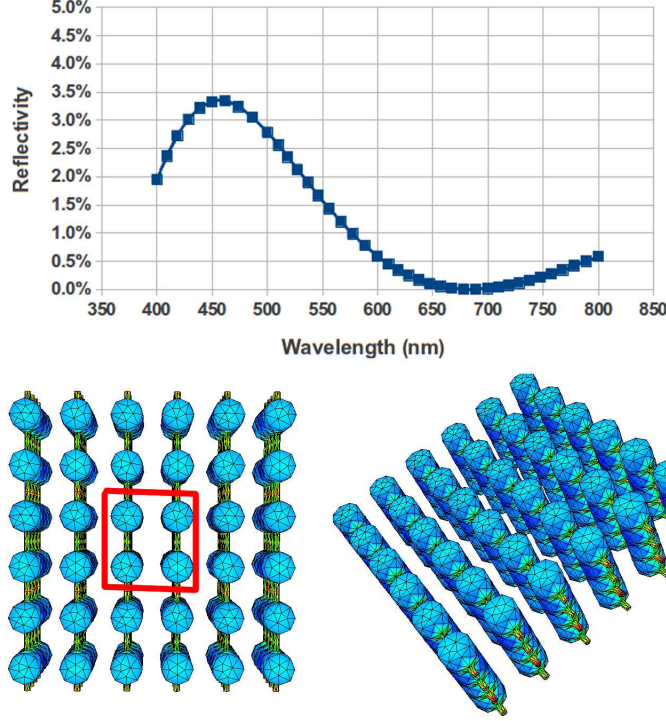


Fig. 8. (Top) Calculated reflectivity of a single model butterfly scale. (Bottom) Scale geometry. A single unit cell with $|\vec{a}_1| = |\vec{a}_2| = 320$ nm is outlined in red. The height of the structure out of plane is 280 nm and the diameter of the larger cylinders is 130nm, with a center-center spacing of 160 nm between nearest neighbors. The smaller cylinders of diameter 20 nm with axes in the plane of periodicity are oriented along the polarization vector of the incident field, with a center-center spacing of 70 nm out of the plane of periodicity. The resultant mesh has $N = 10,782$. Average time to solution (without acceleration, extrapolated): ~ 1430 minutes. Average time to solution (ACE acceleration): ~ 31 minutes. Total speedup: $\sim 46\times$.

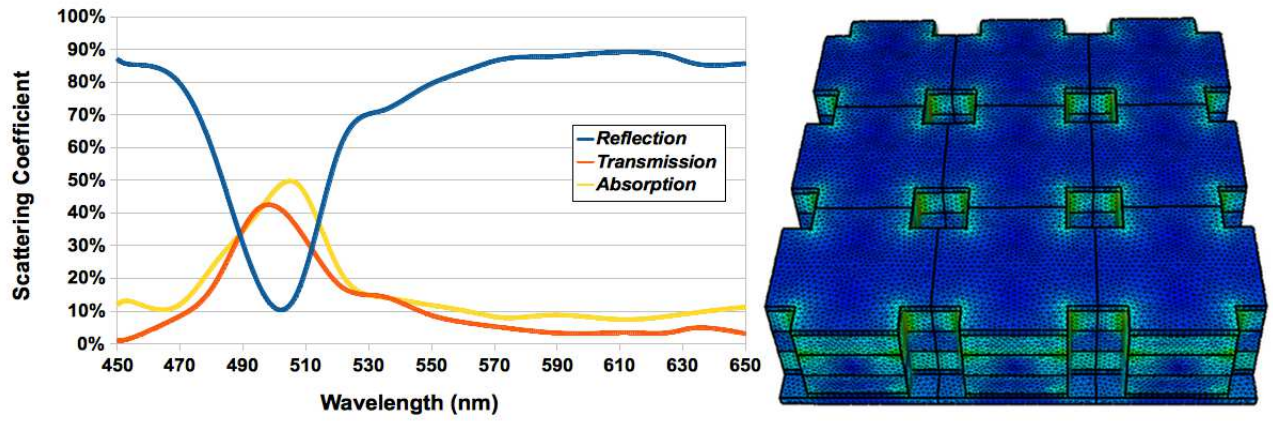


Fig. 9. Demonstration of capability in solving a large scattering problem ($N = 147,374$) inspired by a metamaterial design presented in [39]. An average of 188 iterations per frequency was required, with an average time per iteration of 3.18 minutes, and an average total solution time of 896 minutes.